Philosophical Issues, 17, The Metaphysics of Epistemology, 2007

THE METAPHYSICAL STATUS OF KNOWLEDGE[1]

Hilary Kornblith
University of Massachusetts, Amherst

Much of the work on naturalistic epistemology has focused on methodological issues, arguing for and exploiting the value of empirical work for an understanding of knowledge. But naturalism involves, of course, not only a commitment to certain methodological views (according to which proper method in epistemology, and philosophy generally, is continuous with empirical science); it involves a host of metaphysical commitments as well. For the naturalistic epistemologist, knowledge is properly viewed as a natural phenomenon. In this paper, I examine just what this commitment amounts to and why one might wish to endorse or reject it.

Questions about the metaphysical status of knowledge are sometimes connected to questions about its scope, in particular, whether human beings alone have genuine knowledge, or whether, instead, non-human animals as well count as genuine knowers. Thus, for example, John Haugeland remarks,

> . . . there are senses in which many animals can be said to "know", "want", and even "understand" things around them; but these are not the same as the senses in which people can be said to know, want, and understand things. Classing these uncritically together is as great an obstacle to insight as was classing whales with fish or the sun with the planets.[2]

Now, here, Haugeland makes clear his commitment to the view that knowledge is properly ascribed to human beings alone, but the question of the scope of knowledge might seem to be quite independent of the question of its metaphysical status. After all, one might certainly view knowledge as a natural phenomenon, and, with perfect consistency, hold that the scope of that phenomenon includes human subjects, but not any other animals. Indeed, when Haugeland compares the issue of the scope of knowledge with the question of whether whales are fish, or whether the sun is a planet, he certainly seems to suggest this very view. After all, the discovery that

whales are not fish was an important scientific discovery. Fish constitute a natural kind, and, as it turns out, although whales bear some superficial similarities to them, they are, in the end, fundamentally different. There is no well-motivated scientific taxonomy which places whales and fish in the same category. And the same may be said about the issue of whether the sun is a planet. These two issues are, quite clearly, concerned with the scope of certain natural phenomena, and if the question about the scope of knowledge is just like them, then it might seem that Haugeland regards knowledge as a natural phenomenon as well.

But this is not Haugeland's view at all, as he makes quite clear. The difference between human knowledge and the kinds of states which non-human animals are capable of turns on a property which Haugeland calls "existential commitment."

> It is the capacity for this sort of commitment that I am inclined to think is relatively recent—almost certainly more recent than language, and perhaps more recent than cities and writing. Like city-building and writing, the possibility of existential commitment is part of a cultural heritage (not just a biological or "natural" capacity).[3]

So on Haugeland's view, the question of whether non-humans have knowledge turns out to be quite different from the question of whether whales are fishes. The latter is a question about a certain natural kind, but the former, it turns out, is not. The property of being a fish is a natural property, and questions about its scope are resolved by scientific investigation. But the capacity to have knowledge is not, Haugeland tells us, a natural capacity; it is, instead, "part of a cultural heritage." It is a mistake, on Haugeland's view, to regard knowledge as a natural phenomenon.

Haugeland is not alone in holding this constellation of views. Thus, Michael Williams states quite flatly, "Knowledge is not a natural phenomenon."[4] Like Haugeland, Williams sees the issue of whether knowledge is a such a phenomenon as tied to the question of whether knowledge is the exclusive property of human beings—a view he endorses—or whether, instead, non-human animals are capable of having knowledge as well. And, like Haugeland, Donald Davidson[5] and Robert Brandom,[6] Williams sees these issues as connected with the question of whether non-human animals are capable of having beliefs about their own beliefs, and making the kinds of self-assessments that require such second-order states. Knowledge, on the view which these philosophers share, requires a kind of self-assessment of which non-human animals are not capable, and this, in turn, they believe, bears directly on the issue of whether knowledge should be viewed as a natural phenomenon.

It is not at all obvious why these views should be seen as interdependent. In this paper, I attempt to sort these issues out. Just what does it mean to claim, or to deny, that knowledge is a natural phenomenon? What relation, if any, does this claim have to questions about whether non-human animals

have knowledge? And what does the possibility of self-assessment, and the capacity for second-order belief, have to do with the question of whether knowledge should be viewed as a natural kind?

I

In order to understand the relationship among these issues, we will need to begin by looking at them separately. And before we can address the question of whether non-human animals have knowledge, we must first address the question of whether non-human animals even have beliefs. Let us start by describing three kinds of individuals.

First, consider the case of plants, which respond differentially to their environment: when sunlight is present, they turn toward it. Sunlight affects the plants in a way which brings about this motion, and it brings it about by bringing about various changes inside the plant. Internal states of the plant are thus reliably correlated with the presence of sunshine, and we may thus speak of these states as "information-bearing states," but they are information-bearing in the very same sense that the internal states of my thermostat bear information about the temperature in my living room or the color of a strip of litmus paper bears information about the acidity of a sample of liquid. These states are stimulus-bound: individuals whose states are information-bearing in only this sense are predictable by way of stimulus-response connections alone. The frog's response to flies is like this.[7] If one rolls a BB in front of the frog, it will suck it up as if it were a fly. Well, we all make mistakes. But the frog will do it a second time if a second BB is rolled in front of it; and a third; and a fourth; and so on. The frog will not learn from its experience. Receptors on the frog's retina are directly wired to states which produce the tongue-flicking response. The state which bears the information that there is a fly in front of the frog is not informationally integrated with other information the frog possesses. Individuals whose information bearing states are all of this sort are exceptionally crude: they include various simple mechanical devices, plants and such animals as paramecia.

But many animals are not like this. Their behavior is not explicable by way of simple stimulus-response connections. Thus, consider the piping plover, a bird which feigns a broken wing when a potential predator approaches its nest.[8] The bird moves away from the nest, thereby distracting the predator, and once the predator is far enough away from the nest, the bird flies off. One might reasonably ask whether this behavior is simply stimulus-bound. After all, there are behaviors like this in the animal world which are simply wired in, like the frog's tongue-flicking, and show no sign of sensitivity to other information which the animal may, in some sense, possess. But the piping plover's broken wing display is not like this. Thus, for example, if a human being walks on a path which would come close to the nest, the plover will approach the individual and perform the broken wing display. But if

that person continues on her path without disturbing the nest, and follows this same path on a daily basis, then the plover will stop performing the display. When a different individual approaches, however, on the very same path, the plover goes back to its broken wing performance. This is not some stimulus-bound behavior. In order to explain it, we need to see the plover as picking up information about its environment and integrating it with stored information about the past. Here we have something requiring more than mere S-R connections.

Finally, there are creatures still more sophisticated than piping plovers. These creatures have states which not only register information about the environment—like thermostats and paramecia—and not only have states which form part of a representational system which interacts in ways dependent on the content of those states—like piping plovers—but these creatures also have second-order states, ones which carry information about the states which carry information about the environment. These second-order states are informationally integrated with the first-order states to form a far more flexible and responsive representational system. Human beings are creatures of this sort.

Where should we draw the line separating individuals who do have beliefs from those who do not? Now most of us would agree that individuals of the first sort described above—plants, thermostats, and so on—have no beliefs. And it is equally clear that individuals of the third sort, human beings, do have beliefs. But where do those in the second category belong? Should they be counted among the believers, along with human beings, or are they merely individuals with information-bearing states—like plants and thermostats—but utterly lacking in beliefs? I have argued[9] that many non-human animals do have genuine beliefs. Brandom, Davidson, Haugeland, and Williams[10] disagree. This is not merely a terminological dispute. It is, instead, a dispute about fundamental issues of both epistemology and metaphysics.

What is at stake here? According to Williams, "The fundamental issue is whether being a believer requires a measure of rational self-awareness: the capacity to assess beliefs for their epistemic credentials." [201] As we have seen, having beliefs is more than just a matter of having information-bearing states. After all, a thermostat, and even a piece of litmus paper, have information bearing states. In order for an information-bearing state to qualify as a belief, according to Williams et al., it must be sensitive to reasons. The color of the litmus paper, for example, is sensitive to features of the external environment, but this is not enough, on this view, to constitute a sensitivity to reasons. As Williams explains,

> What we [Williams, Davidson and Brandom[11]] argue, in our different ways, is that sensitivity to reasons or epistemic responsibility requires the *capacity* for assessing the credibility of one's commitments, revising them if need be. This capacity requires extensive knowledge of one's epistemic strengths and weaknesses: for example, the sorts of things one is or is not good at recognizing. [208]

Sensitivity to reasons, then, is a matter of epistemic responsibility which, in turn, requires a capacity for certain sorts of self-knowledge. Thus, on this view, it is impossible to have beliefs unless one also has second-order beliefs; it is impossible to have knowledge without also having second-order knowledge. To put this in the form of a slogan: No cognition without meta-cognition. Williams explicitly endorses this way of viewing things. "... the main point ... has to do with whether knowledge of the environment is essentially bound up with ... 'meta-cognitive processing'. Brandom, Davidson and I all think that it is." [194]

Now one might think that a natural view to hold about many non-human animals is that they have first-order beliefs—beliefs about their environment—but that they have no second-order beliefs at all. Many non-human animals are thus believers; they are cognitive beings, even though they are incapable of meta-cognition. This is the view I favor. But if Williams et al. are correct, the class of beings who have first-order beliefs and yet are utterly lacking in second-order beliefs is, of necessity, empty. It is only possible to have first-order beliefs if one also engages in a substantial amount of self-scrutiny, evaluating one's own cognitive position. One certainly needn't evaluate the epistemic credentials of every belief one holds; absolutely no one does this, and, indeed, no one could. But one simply doesn't count as a cognitive being—a creature with beliefs and other propositional attitudes—unless one does, at least at times, engage in a certain amount of self-scrutiny. Without such self-scrutiny, one does not count as forming beliefs; one is merely, like the piece of litmus paper, registering information. This kind of self-scrutiny, however, requires beliefs about one's own beliefs. If one is to take responsibility for one's beliefs, one must direct one's attention to them and, at least at times, consider the possibility that one might be wrong. This, of course, requires forming beliefs about one's beliefs. One thus cannot even have first-order beliefs unless one has second-order beliefs as well. Thus, anyone who argues that non-human animals lack second-order beliefs, as I have,[12] thereby succeeds, as Williams would have it, in making the case that such animals lack beliefs of any kind.

Thus, as Williams et al. see it, having beliefs requires being sensitive to reasons; being sensitive to reasons requires epistemic responsibility; epistemic responsibility requires attentiveness to one's capacities, strengths and weaknesses, abilities and limitations; and this, in turn, requires beliefs about one's own beliefs. On this view, then, there are individuals with informational states but no beliefs; and there are creatures with both first-order as well as second-order beliefs; but there simply could not be any creatures with first-order beliefs but no second-order beliefs at all.

II

How do things look from the other side? Let us return to our three categories of individuals, representative members of which included plants, plovers

and people. As I see it, the striking difference between the first and the second category is this: while the first category includes individuals with information-bearing states, the states in these individuals do not add up to a representational system; in the second category, however, the information-bearing states do constitute such a system. As has been pointed out, it is not only plants which have information-bearing states; such states are also possessed by thermostats and even pieces of litmus paper. The fact that an individual has information-bearing states, by itself, does not show that the individual has any sort of psychology at all. When we move, however, from organisms with states which bear information about the environment to those with states which form a representational system, we thereby make the transition to creatures with propositional attitudes. As I see it, these are creatures with full-blown beliefs, desires and so on; these are genuine cognitive beings.

When we look at the plover's sophisticated behavior, we see not only that it registers information about its environment, but these internal information-bearing states interact with one another to allow for more sensitive interactions with the environment. The plover first registers information about a potential threat—the human beings walking toward the nest—and takes appropriate action by engaging in the broken wing display. When the plover discovers, after repeated interactions with this particular individual, that no threat is present, she ceases to respond by engaging in the display. This requires the integration of new information with old, and this is why we need to posit a representational system to account for such sophisticated behavior. But, as I see it, there is now every reason to regard these informational interactions as cases of inference: they are, after all, transitions involving the interaction among representational states on the basis of their content.

Consider a second case of sophisticated non-human behavior, the case of tool use in apes. Here is Richard Byrne's description of a famous such case:

> The wild chimpanzee is "presented" with a termite mound, initially with no hole, in the African bush. (To be precise, only a few dozen rather inconspicuous termite nests are located in the home range of 15-20 square kilometers of forest and bush. A random day's walk would miss all of them.) The range of objects with which a mound might be probed (or prodded, or struck) is huge; very few of these natural objects are tried. Even the hole has to be made, and this can only be done in certain parts of the mound where termites are getting ready to emerge, at a certain time of year... The edible termites are not visible at all. Probing must be done delicately, and often the other hand is used as a guide and support. Hasty poking destroys the tool; hasty pulling out dislodges any attached termites. There is, in short, a lot to learn: termite fishing is a sophisticated skill.[13]

The chimpanzee, it seems to me, understands what it is doing. The sophistication of the behavior, the way in which the various subtle details of

termite fishing are learned, the way in which the chimpanzee is responsive to success and failure of prior attempts, all tell in favor of viewing this behavior as involving significant cognitive achievements. Williams et al. frequently emphasize the connection between having beliefs and understanding, and I agree that there is such a connection. Plants, thermostats and pieces of litmus paper have no understanding of their environment. But it is clear from the above description that chimpanzees do have such understanding. While they may be utterly lacking in an understanding of their own mental states, such a lack of understanding does not in any way preclude an understanding of the external environment.

I also want to agree with Williams et al. that having beliefs requires a sensitivity to reasons, and I believe that the behavior of the plovers and the chimpanzees shows that they are sensitive to reason. When the plover encounters a reason to stop engaging in the broken wing display, it stops; when the chimpanzee encounters a reason to modify its method of fishing for termites, it responds appropriately. As I see it, meta-cognitive evaluation is only one way to show a sensitivity to reasons. Just as constant self-monitoring is unnecessary for such sensitivity, self-monitoring of any sort is unnecessary. This is not to say that one might not achieve a greater degree of sensitivity by way of meta-cognitive monitoring. But, at the same time, one should not simply assume that greater reliability is inevitably the product of some such self-scrutiny.[14] Once one sees that meta-cognitive evaluation may sometimes result in less sensitivity to reason, sometimes result in no change in one's sensitivity to reason, and sometimes result in greater sensitivity to reason, one is faced with a choice: does one care about such self-scrutiny regardless of its efficacy in improving reliability, or does one care about it only insofar as it is instrumental in producing a greater sensitivity to reason? I don't think that the first of these two options is reasonable. But the second option makes reliability fundamental, and views meta-cognition as one process among many which may, in some cases, aid in generating a responsiveness to reasons.

Consider how we are to tell the story of cognitive development in children, that is, creatures who make the transition from the second of our categories to the third. As developmental psychologists tell us,[15] mental state concepts are fairly late arrivals on the conceptual scene: children do not have beliefs about beliefs, for example, until roughly the age of four. Now one way to tell the story of cognitive development, indeed, the way most developmental psychologists tell the story, is that, early on, children have a robust set of beliefs about their physical environment; it is only quite a bit later that they come to have beliefs about mental states. If such an account is correct, as I believe it to be, then we may give the obvious account of the sophisticated behaviors of three year olds, namely, that they have an extensive set of beliefs about their environment, and elaborate desires as well, and these beliefs and desires jointly conspire to produce behavior. Children

under the age of four learn a great deal about their environment, and the natural account of this learning requires that they be seen as engaging in inference: new information is not simply registered; it is combined with pre-existing knowledge to produce new knowledge. If we see things in this way, as I would suggest we should, then children under the age of four will be regarded as having propositional attitudes, but they will not be seen, from the beginning, as having second-order propositional attitudes. The various concepts of the propositional attitudes themselves—the concepts of belief and desire, for example—are far more complex than the concepts acquired early on. Needless to say, the concepts of truth and falsity are extraordinarily sophisticated; it should be no surprise that various physical object concepts are acquired long before there is any understanding of the idea of truth.

If we insist, however, that it is impossible to have first-order beliefs unless one also has second-order beliefs, then the developmental story becomes far more difficult to tell. One will need an explanation of the kinds of learning which occur before the age of four which does not in any way appeal to beliefs or desires. On this account, if we acknowledge that children do not acquire the concept of belief until age four, we will need to insist that there are no beliefs of any kind before that age, and no desires, presumably, as well. At the same time, it must be recognized that these children are not simple S-R individuals; they are far more sophisticated than plants and thermostats. One may, of course, insist that the term 'belief' is an honorific which is to be reserved only for creatures who have second-order beliefs as well as first-order states, but the motivation for this kind of insistence is beginning to look strained. The explanation of early behavior would now require "proto-beliefs" and "proto-desires." Only when these states themselves become mentally represented, as part of the very same representational system, do they attain the status of genuine beliefs and desires. There seems to be nothing in the phenomena themselves which would justify such a terminological distinction. More than this, one is forced to tell a story about concept acquisition which is extremely implausible: instead of supposing that physical object concepts are acquired early on and that mental state concepts are later arrivals, as the standard story would have it, Williams et al. are forced to say that there are no physical object concepts prior to the acquisition of mental state concepts. It is hard to make sense of such an account of cognitive development.

Note, in particular, that the standard proposed by Williams and Brandom—no beliefs without meta-beliefs—is far more demanding than the suggestion that there are no beliefs without language.[16] Children not only have a very substantial vocabulary at age 3; they are speaking in complicated sentences. Thus, Steven Pinker remarks,

> Between the late 2s and mid-3s, children's language blooms into fluent grammatical conversation so rapidly that it overwhelms the researchers who study it; no one has worked out the exact sequence. Sentence length increases steadily

> and, because grammar is a combinatorial system, the number of syntactic types increases exponentially, doubling every month, reaching the thousands before the third birthday.[17]

Thus, if we follow Williams and Brandom, then it is not just piping plovers and chimpanzees who will be resigned to the realm of non-believers; it is children below the age of four as well, in spite of the fact that many of them are fluent language users. This is, to say the least, hard to credit. Davidson is no better off. While he is well-known for his view that there are no beliefs without language, this is not to say that he thinks language use is sufficient for belief. Indeed, what he suggests is that one cannot have beliefs unless one has "the idea of an objective, public truth."[18] It is safe to say that such a concept is not possessed by four-year-olds. Haugeland, as the quotation at the beginning of this paper reveals, is less generous still. Thus, in all of these authors, the standard account of cognitive development must be rejected in favor of an account which sees proto-beliefs giving way to beliefs when some critical juncture—the acquisition of second-order beliefs; the acquisition of the concept of truth; the cultural acquisition of "existential commitment"—is reached. Any such account of cognitive development is terribly implausible.

III

I have been arguing that the standard account of belief which we find in developmental psychologists runs contrary to the suggestions of Brandom, Davidson, Haugeland and Williams, and that these philosophers offer us an account of belief which will make it difficult to tell a plausible story about cognitive development. But perhaps this kind of argument begs the question against these authors by presupposing a claim they would all reject: that belief is a natural kind. I have, after all, supposed that there is a certain natural phenomenon, believing, which occurs in human adults, and I have been asking whether this very same phenomenon occurs in non-humans and young children. I have assumed, as a working hypothesis, that there is indeed a unified phenomenon of believing, and that it is possible to give an account of it which explains the causal role which beliefs play in producing behavior. Developmental psychologists in fact offer just such a unified explanatory account, and those who set significantly higher standards for belief thereby offer a view of this phenomenon which stands in the way of a plausible unified account. To my mind, this counts in favor of the view of the developmentalists, and against the views of Brandom, Davidson, Haugeland and Williams.

But perhaps this is unfair. Haugeland explicitly rejects the view that the propositional attitudes constitute a natural kind. He says that they are "part of a cultural heritage (not just a biological or "natural" capacity)."[19] And

Williams says that "knowledge is not a natural phenomenon";[20] he makes quite clear that he has the same view of belief, and that he sees this view in Brandom and Davidson as well. So if these philosophers reject the view that belief is a natural kind, that the notion of belief they are trying to articulate must therefore fit into some well-confirmed scientific theory which explains the causal role which beliefs play in producing behavior, then one might reasonably think it unfair to argue against their views by pointing out that they fail to fit in with any such account.

Let me therefore acknowledge that these philosophers are not trying to offer a scientific account of belief, or even an account of belief which is compatible with the best available scientific theories. Even if there is a natural kind of which the developmental psychologists are giving a well-confirmed account, that natural kind is not the object of study which Brandom, Davidson, Haugeland and Williams are talking about. They are talking about something else.

So what is it that these philosophers mean to be talking about, if it is not the natural phenomenon which developmental psychologists study? Merely rejecting the view that belief is a natural kind does not, of course, relieve one of the responsibility to locate the subject of one's investigation. These philosophers attempt to offer an account of the nature of belief; they merely deny that belief, as they understand it, is an object of scientific study.

Thus, for example, if a marine biologist attempts to explain what sponges are, and someone who is interested in household products should come along and insist that he too is interested in sponges, but not the natural kind that goes by that name, I think none of us would have any trouble understanding what is being claimed. There is a natural kind which marine biologists refer to with the word 'sponge.' There is a household cleaning product which is also referred to as a 'sponge.' These are two different things, and there is no need for the theory of the household product to answer to the science of marine biology. Household sponges are not a natural kind, one might reasonably say. The word 'sponge' is simply ambiguous.

So unless these philosophers wish to reject the scientific account of belief which is embedded in current work in developmental psychology—and none of them attempt to engage with that literature—then we should assume that they view any such account as no more relevant to what they are doing than the marine biological account of sponges is to the investigation of household products. The latter investigation may have an integrity of its own without taking on the (irrelevant) work of marine biologists.

If this is the right way to see what those who reject the view that belief and knowledge are natural kinds, then some case needs to be made that there is, as in the case of sponges, a genuine ambiguity in our talk of beliefs, with a non-scientific notion sitting side by side with the scientific one. And we need to have some sense, as we clearly do in the case of sponges, of what the phenomenon is which the philosophical account is supposed to answer to.

In the case of belief, however, this task is complicated by the fact that the non-scientific notion of belief which these philosophers mean to explicate seems to be called on to explain many of the very same phenomena which the scientific notion addresses. Thus, all of the philosophers under discussion view beliefs as causes of behavior. In addition, explanations are sometimes offered which seem to be presented as rivals to the scientific account. For example, Michael Williams defends the implication of his account that non-human animals do not have beliefs by attempting to locate an important difference between human and non-human mental representations:

> We might say, animals don't need the capacity for epistemic assessment because they don't test hypotheses: they test themselves. But this is why they are not truly sensitive to reasons. They cannot really change their minds, though the information-acquiring and processing capacities of the species can change over time. [209]

This kind of argument puzzles me. Williams seems to be arguing that there is an important difference between human belief and the mental representations of non-human animals, and because they are different in this fundamental way, they should not be seen as belonging to the same kind. While I am certainly sympathetic in principle with this sort of argument, this sounds like an argument that animal mental representations and human beliefs constitute different natural kinds; it doesn't sound at all like an argument for the view that human belief doesn't constitute a natural kind at all.

But this is not the only problem with Williams' argument, for the claim he makes about the difference between human and non-human animals is also untrue. Williams presents a picture of non-human animals as acquiring information about their environment by way of processes which are hard-wired and thus insensitive to new information. According to Williams, "the information-acquiring and processing capacities of the species can change over time" in the case of non-human animals, but, unlike humans, they do not change within the life of an individual. This is certainly the case with the frog's hard-wired response to flies and BBs. It might change in the species in response to evolutionary pressures, if, for example, frogs were to find themselves in an environment where moving BBs are plentiful, but this process of information acquisition is not going to change in the life of a single frog; it does not respond to new information the frog comes to possess. But this is not true of animal cognition generally. As we have seen, the plover treats information about those approaching the nest differentially, depending on its past experience with those very individuals. And the ways in which chimpanzees approach a termite mound depends on their prior successes and failure in the enterprise of termite fishing. Nor are these cases rare in animal cognition. They are extraordinarily common.[21] How information is

processed does, in fact, change over time within the life of individual animals; just as in human beings, when new information is acquired, it may change the way in which subsequent information is processed.

Williams' suggestion that non-human animals "don't test hypotheses: they test themselves" is very much in the spirit of Quine's famous suggestion that "creatures inveterately wrong in their inductions have a pathetic but praiseworthy tendency to die before reproducing their kind."[22] Williams sees the tendency of animals to die if their inferences are inaccurate, however, as the only means by which their inferential accuracy may be improved. This presents us with an extremely impoverished view of animal learning, and a view thoroughly at odds with the facts. The suggestion that non-human animals do not have beliefs cannot be justified in this way.

If we wish to understand the case against viewing belief and knowledge as natural kinds, we will thus need to examine other considerations.

IV

Williams et al. hold the view that in order to have beliefs of any kind, a creature must also have second-order beliefs. Without beliefs about one's own beliefs, a creature is not, on this view, capable of epistemic responsibility; it is not, therefore, sensitive to reasons. At most, it is capable of responding to changes in its environment. Such sensitivity to environmental features, however, is not at all the same as sensitivity to reasons, as the litmus paper example vividly illustrates.

Thus far, I have been presenting certain external difficulties which this sort of view faces. But there is a problem internal to the view as well, a problem which, I believe, completely undermines it. In this section of the paper, I present this internal difficulty. I then return to the issue of whether belief and knowledge are natural kinds.

Consider, once again, the piping plover. When an individual approaches the nest, this produces in it an information-bearing state which represents the individual as a threat. But when this individual has approached the nest repeatedly without harming the nest, the plover no longer represents the approach of the individual as threatening. The plover represents and stores information about the repeated harmless approaches of this particular individual; these information-bearing states causally interact with the state which bears the information that the individual is approaching again, and this, in turn, produces the information-bearing state that the individual is not a threat. I see these information-bearing states as beliefs, and the causal interaction among them as inferential; I regard the information that previous encounters with this individual were harmless as providing the plover with a reason for forming the belief that this next encounter will be harmless as well.

But Williams et al. see this differently. The plover has, as I would agree, no second-order beliefs. So, even on my account, the plover is not in a position to examine its first-order beliefs, to consider the question of whether they constitute good reason to form the belief that the individual approaching the nest is harmless. Without such second-order beliefs, Williams argues, the first-order states which register information about the past behavior of the individual approaching the nest do not constitute reasons for belief, and since they do not constitute reasons, the plover is not sensitive to reasons. Because it is not sensitive to reasons, it should not be seen as having even first-order beliefs.

So now consider a more sophisticated individual, a human being, a creature who has both first- and second-order beliefs. Let us suppose that this person, let's call him Jack, believes that p, and also believes that if p then q, and forms the belief that q on the basis of these two beliefs. Let us further suppose that before Jack formed the belief that q, he examined his first-order beliefs and noted that he believed that p and that if p then q; that is, he formed the second-order belief that he believed both p and if p then q. Only after noting this did Jack go on to believe that q. Now it is precisely Jack's capacity to engage in this sort of thinking that makes Williams et al. regard him as a creature who is sensitive to reasons, unlike the plover. What role did Jack's second-order belief play in his reasoning? Now one possibility is that it played no role at all; it was causally inert. Jack simply formed this belief about his first-order beliefs and then, for reasons entirely unconnected with the second-order belief, went on to form the belief that q. Now if this were the role which second-order beliefs were thought to play in reasoning, namely no role at all, it would be hard to see why we should think that Jack counts as any more sensitive to reason than the plover. If Jack's belief that q is simply caused by his beliefs that p and if p then q, with his second-order beliefs playing no role at all, then surely Jack's second-order belief can do no epistemic work for Jack. And if the plover is thought not to count as having beliefs, it would be hard to see why forming second-order beliefs in this way, beliefs which play no role in one's reasoning, should mark the difference between creatures wholly lacking in beliefs, such as the plover, and creatures who do have beliefs, such as humans.

So Williams et al. must think that what makes the difference here is not just forming the second-order belief, but having this belief play a causal role in one's reasoning. The idea seems to be that the beliefs that p and that if p then q cannot constitute reasons for Jack to believe that q unless he also forms the second-order belief that he believes both p and if p then q.[23] But if these first-order beliefs do not constitute reasons for forming the belief that q unless there is some higher-order belief about them, then we may ask the same question about that higher-order belief: Does it constitute a reason for believing that q? Suppose that it, together with the beliefs that p and if p then q jointly cause the belief that q. Does Jack then form the belief that

q on the basis of reasons? Williams has already told us that the first-order beliefs do not count as reasons unless they are backed by a higher-order belief. But now, it seems, he must say the same about the second-order belief. If it simply causes some information-bearing state, without Jack recognizing that he has that second-order belief, then it does not constitute a reason for believing. In order for it to constitute a reason for Jack to believe, Jack would have to form a belief about it. But now we are launched on an infinite regress.[24]

The problem is straightforward. If we insist that first-order beliefs do not constitute reasons for belief unless they are backed by higher-order beliefs about them, then we should hold that second-order beliefs do not constitute reasons for belief unless they are backed by higher-order beliefs about them. This requires an infinite hierarchy of beliefs if one is to have any beliefs at all. It will not only eliminate non-human animals from the ranks of believers; it will eliminate human beings from the ranks of believers as well.

This approach to the question of what standards need to be met in order for an individual to count as a believer thus sets not only an unreachable standard for belief; it also sets an unreasonable standard for belief. Having beliefs just isn't this difficult.

V

Let us return then to the issue of whether belief and knowledge constitute natural kinds. Leaving aside the problem which is generated by the requirement that higher-order states play a causal role if reasoning is to take place, why should we think that neither belief nor knowledge constitute real kinds in nature?

When Haugeland addresses this issue, he contrasts the idea that propositional attitudes are biological or natural kinds with the suggestion that they are "part of a cultural heritage."[25] Williams too brings up the issue of a social dimension to propositional attitudes in this context.

> Consideration of the demands of epistemic responsibility helps to clarify the ways in which knowledge is and is not a social undertaking. To become sensitive to reasons is to learn what can and cannot reasonably be questioned, and why. Thus acquiring sensitivity to reasons requires a lot of know-how, as well as extensive meta-knowledge. To become a knower is to become acculturated. This is not something that one could, even in principle, do for oneself. [208]

Brandom and Davidson, each in their own way, also argue that there are certain social dimensions to the propositional attitudes. These suggestions raise two different questions: First, are there important social dimensions to the propositional attitudes? And second, if the propositional attitudes do, in

some sense, have a social dimension to them, does this thereby show that they do not constitute natural kinds? I will address each of these questions in turn.

Let us look at the quotation from Williams carefully. "To be sensitive to reasons," he tells us, "is to learn what can and cannot reasonably be questioned, and why." Much of human learning certainly involves learning "what can and cannot reasonably be questioned," and much of this is straightforwardly social. As Williams remarks, acculturation involves this very sort of learning, and to learn what can and cannot reasonably be questioned certainly does allow one to become sensitive to reasons. So this sort of cultural process is one way in which one might become sensitive to reason. Talk of "questioning" here may also suggest something typical of the human case, namely language use, something crucial, of course, in Davidson's account. But social learning, by itself, needn't involve language use. Chimpanzees learn how to fish for termites from their conspecifics, but the skill of termite fishing is not taught in a way which involves linguistic communication. One might emphasize social dimensions to learning without thereby committing oneself to any view at all about the essentiality of linguistic communication, even if such communication is, of course, the norm in human social learning.

Talk of "what can . . . be questioned," however, while it certainly suggests something linguistic, and social as well, need not perhaps involve either of these features. I may, without any social interaction at all, question myself about what I ought to believe. And if all one means by coming "to learn what can and cannot be reasonably questioned," is coming to learn how inquiry ought to proceed—which possibilities to take seriously and which one may reasonably ignore—then the linguistic dimension seems optional here as well. The chimpanzee, in learning to fish for termites, learns to take certain things for granted and to carefully check on others. And although the chimpanzee learns this particular skill socially, not everything learned is learned in this way. Interaction with the world is one good way to learn which things one may take for granted, and which things need to be checked. The world, after all, has a way of informing one if one fails to check on things which should not be taken for granted.

Now the kind of learning "what can and cannot reasonably be questioned" which is essential for having a sensitivity to reasons is surely the kind involved in coming to learn how inquiry ought to proceed. But it is now clear that, even if this learning is typically both social and linguistic in the human case, such learning requires neither social nor linguistic contributions. Given sufficiently flexible native capacities, interaction with the world may teach one how to proceed with one's inquiries by teaching one what one should test and what one may simply rely upon. Creatures who have such sophisticated capacities, whether they have second-order beliefs or not, whether they have linguistic capacities or not, and whether their learning is social or not, thereby

display a sensitivity to reasons. The suggestion that a sensitivity to reasons is essentially social, or cultural, is thus simply mistaken.

But let us suppose, just for the sake of argument, that sensitivity to reasons was essentially social or cultural. Why should one think, as Williams and Haugeland clearly do, that this would show that belief and knowledge are not natural kinds? Why should we think that "social" and "natural" are contraries? Here is what Williams says on this score, in explaining why he is opposed to the idea that epistemology should be naturalized.

> . . . to claim that knowledge need not be naturalized does not mean that it should be understood in some "anti-naturalist" (let alone supernaturalist) way. It is to claim, rather, that although we can illuminate the concept of knowledge, and explain why it matters to us to have a concept like that, knowledge is not a proper object of theory, scientific or otherwise. [209]

So epistemology, on Williams' conception, will "illuminate the concept of knowledge," and "explain why it matters to us to have a concept like that;" but, at the same time, we should think that "knowledge is not a proper object of theory." What does this mean?

Consider, by way of contrast, the concept of disease. Historians of medicine might attempt to illuminate the concept of disease and explain how it has changed over time. They might, in addition, explain why it matters to us to have a concept like that. But in this case, of course, we do have a proper object of scientific theorizing, and it is precisely by way of successful scientific theorizing about the nature of disease that we may shed light on why it is so important for us to have a concept of such a thing. It is only by understanding what diseases really are that we may understand why it is important for us to have a concept of them. So in this case, the kinds of conceptual investigations which Williams favors in the case of knowledge simply are not at odds with the kind of naturalistic investigation I endorse. But in the case of disease, although the concept is, in some sense, socially constructed, and our knowledge of diseases is socially mediated, disease itself is not a social kind, and it is for this reason, no doubt, that Williams would see the case of disease as different in kind from the case of knowledge.

So let us look at an example of a social kind. Consider the case of "democratically elected government." This is a kind which is social through and through if anything is. Here we may certainly attempt to illuminate the concept of a democratically elected government, trying to say just what this kind of government comes to; and we may attempt to say something about why it is important that we have a concept of this kind. But here too, it seems, we may theorize about democratically elected governments themselves, and attempt to discover what the connection is, if any, between this form of government and other social conditions. Indeed, this very form of theorizing is familiar in the social sciences. The fact that this is a socially constructed

kind does not in any way make it an improper object of theorizing. Social kinds may, after all, have a degree of stability that allows for significant causal explanations in terms of them.[26] Indeed, it is for this very reason that social theories may be so illuminating. If the social kinds were not an object of proper theorizing, then the very idea of social science would be a sham. I take it that Williams is not endorsing the radical claim that there is no such thing as a legitimate social science.

But why then should we think that there is some essential conflict between the alleged social character of knowledge, or our concept of knowledge, and regarding knowledge as a proper object of theory? The social dimension of knowledge is important for Williams, I believe, not simply qua social, but rather because, on his view, it provides the grounding for normativity. In the end, it is not that the social and the natural are seen as automatically opposed by Williams, but rather that the normative and the natural are seen as antithetical. Thus, in summarizing Brandom's view on the nature of reasons, a view with which he has a good deal of sympathy, Williams remarks,

> . . . we cannot capture content-determining inferential relations without recourse to *normative* vocabulary. Reasons must be intelligible as *good* reasons. Moreover, those to whom we attribute sensitivity to reasons must be such that we can see them as taking *normative attitudes*—treating some reasons as good, others as not so good—for this is what, in practical terms, sensitivity to reasons comes to. [201]

As a number of philosophers now put it, what is at issue here is whether there is a fundamental split between the "space of reasons" and the "space of causes."

We can agree with Williams (and Brandom) that sensitivity to reasons requires "treating some reasons as good, others as not so good." It is one thing, however, to show a sensitivity to reasons, quite another to conceptualize reasons *as* reasons.[27] As we have seen in our earlier discussion of second-order belief, sensitivity to reasons may be achieved even without having the concept of a reason, or of a belief. If all Williams means by his talk of "normative attitudes" is the differential treatment of reasons—treating some as good and others as not so good—then normative attitudes are had by non-human animals, and they are not the exclusive property of human beings. If, on the other hand, talk of normative attitudes is meant to be more demanding, requiring the possession of normative concepts, then we have been given no reason at all to think that sensitivity to reasons requires having attitudes of this kind.

Wherever we stand, however, on the question about what a sensitivity to reasons requires, we have been given no reason at all to think that normative features of belief and knowledge cannot be objects of theory. Work in cognitive ethology and human cognitive psychology deals with the

very sensitivities to which Williams directs our attention. There is important theoretical work being done on the nature of concepts, and how it is that concepts play a role in cognition. Whether these concepts be social or not, whether they be normative or not, has proven simply irrelevant to productive theorizing. In the face of these successful programs of research, it simply will not do to insist that knowledge cannot be an object of theory.

VI

We have examined several different avenues for defending the view that belief and knowledge cannot be natural kinds: the suggestion that belief, and sensitivity to reason, requires second-order mental states; the suggestion that the propositional attitudes are somehow inherently social; the suggestion that propositional attitudes are inherently normative; the suggestion that belief and knowledge are not proper objects of theory. In each of these cases, there seems little reason to think that, to the extent that these ideas have merit, they are opposed to the view that belief and knowledge might be natural kinds. If there is reason to reject the view that belief and knowledge are indeed natural kinds, then it seems unlikely to be found in the constellation of considerations examined in this paper.

Notes

1. A brief version of this paper was presented at an Author Meets Critics session of the Pacific Division of the APA in March, 2004, in response to Michael Williams' comments on my book, *Knowledge and Its Place in Nature*, Oxford University Press, 2002. I especially want to thank Michael Williams for extremely useful discussion on that and many other occasions. I am also indebted to comments from the audience at the APA, and to Kirk Michaelian for helpful written comments on a draft of the APA paper.
2. *Having Thought: Essays in the Metaphysics of Mind,* Harvard University Press, 1998, 2.
3. *Ibid., loc. cit.*
4. "Is Knowledge a Natural Phenomenon?," in R. Schantz, ed., *The Externalist Challenge*, de Gruyter, 2004, 194. Page numbers in the text all refer to this article.
5. See especially "Thought and Talk," in *Inquiries into Truth and Interpretation*, Oxford University Press, 1984, 155–170, and "Rational Animals," in *Subjective, Intersubjective, Objective*, Oxford University Press, 2001, 95–105.
6. See *Making It Explicit: Reasoning, Representing and Discursive Commitment*, Harvard University Press, 1994, and *Articulating Reasons: An Introduction to Inferentialism*, Harvard University Press, 2000.
7. J. Y. Lettvin, H. R. Maturana, W. S. McCulloch and W. H. Pitts, "What the Frog's Eye Tells the Frog's Brain," in Warren McCulloch, *Embodiments of Mind*, MIT Press, 1965, 230–255.

8. Carolyn Ristau, "Aspects of the Cognitive Ethology of an Injury-Feigning Bird, the Piping Plover," in Carolyn Ristau, ed., *Cognitive Ethology: The Minds of Other Animals*, Lawrence Erlbaum Associates, 1991, 91–126.
9. *Knowledge and Its Place in Nature*.
10. In *Knowledge and Its Place in Nature*, I described Williams' view as one according to which non-human animals do have beliefs, but do not have knowledge. It is now entirely clear that this is not Williams' view: he rejects the claim that non-human animals have beliefs. See "Is Knowledge a Natural Phenomenon?"
11. Williams is explicitly discussing these authors in the quoted passage. But the remarks he makes apply equally, it may be argued, to Haugeland as well.
12. *Op. cit.*
13. "Social and Technical Forms of Primate Intelligence," in F. de Waal, ed., *Tree of Origin: What Primate Behavior Can Tell Us about Human Social Evolution*, Harvard University Press, 2001, 163.
14. I have argued this point at length in chapter four of *Knowledge and Its Place in Nature*.
15. See, for example, Janet Wilde Astington, *The Child's Discovery of the Mind*, Harvard University Press, 1993; Alison Gopnik and Andrew Meltzoff, *Words, Thoughts and Theories*, MIT Press, 1997; Janet Astington, Paul Harris and David Olson, eds., *Developing Theories of Mind*, Cambridge University Press, 1988.
16. Similar difficulties face José Bermudez, as I point out in my response to his critical remarks on *Knowledge and its Place and Nature*. See Bermudez, "Knowledge, Naturalism and Cognitive Ethology: Kornblith's *Knowledge and Its Place in Nature*," *Philosophical Studies*, 127(2006), 317–335; Kornblith, "Reply to Bermudez and BonJour," *Philosophical Studies*, 127(2006), 337–349. Bermudez develops his view at length in *Thinking without Words*, Oxford University Press, 2003.
17. "Language Acquisition," in Lila Gleitman and Mark Liberman, eds., *An Invitation to Cognitive Science*, vol. 1, 2nd edition, MIT Press, 1995, 143–4.
18. "Thought and Talk," 170.
19. *Having Thought*, 2.
20. "Is Knowledge a Natural Phenomenon?," 194.
21. See, for example, John Alcock, *Animal Behavior: An Evolutionary Approach*, 8th edition, Sinauer Associates, 2005; Colin Allen and Marc Bekoff, *Species of Mind: The Philosophy and Biology of Cognitive Ethology*, MIT Press, 1997; Sara Shettleworth, *Cognition, Evolution and Behavior*, Oxford University Press, 1998; Jacques Vauclair, *Animal Cognition: An Introduction to Modern Comparative Psychology*, Harvard University Press, 1996.
22. "Natural Kinds," in *Ontological Relativity and Other Essays*, Columbia University Press, 1969, 126.
23. For a similar view, see Sydney Shoemaker, "On Knowing One's Own Mind," reprinted in his *The First Person Perspective and Other Essays*, Cambridge University Press, 1996, 25–49.
24. This problem does not arise for Shoemaker's version of the view since he holds that higher order belief

> is supervenient on first-order beliefs and desires—or rather, it is supervenient on these plus a certain degree of rationality, intelligence and conceptual capacity. By this I mean that one has the former *in* having the latter—that having the former is nothing over and above having the latter. [*op. cit.*, 34]

There are, I believe, problems with this sort of view as well, but they would take us far afield from the subject of this paper. I have discussed this issue in "What Reflective Endorsement Can't Do," manuscript.

25. *Having Thought*, 2.
26. This point is nicely developed in Ron Mallon's "Social Roles, Social Construction and Stability," in Frederick F. Schmitt, ed., *Socializing Metaphysics: The Nature of Social Reality*, Rowman and Littlefield, 2003.
27. Thus, note, someone may show a tremendous sensitivity to ragweed without conceptualizing ragweed *as* ragweed.